The Exact Rate-Memory Tradeoff for Caching with Uncoded Prefetching

Qian Yu*, Mohammad Ali Maddah-Ali†, and A. Salman Avestimehr*
* Department of Electrical Engineering, University of Southern California, Los Angeles, CA, USA
† Nokia Bell Labs, Holmdel, NJ, USA

arXiv:1609.07817v1 [cs.IT] 25 Sep 2016

Abstract

We consider a basic cache network, in which a single server is connected to multiple users via a shared bottleneck link. The server has a database of a set of files (content). Each user has an isolated memory that can be used to cache content in a prefetching phase. In a following delivery phase, each user requests a file from the database and the server needs to deliver users' demands as efficiently as possible by taking into account their cache contents. We focus on an important and commonly used class of prefetching schemes, where the caches are filled with uncoded data. We provide the exact characterization of rate-memory tradeoff for this problem, by deriving the both the *minimum average rate* (for a uniform file popularity) and the *minimum peak-rate* required on the bottleneck link for a given cache size available at each user. In particular, we propose a novel caching scheme, which strictly improves the state of the art by exploiting commonality among users' demands. We then demonstrate the exact optimality of our proposed scheme through a matching converse, by dividing the set of all demands into types, and showing that the placement phase in the proposed caching scheme is universally optimal for all types. Using these techniques, we also fully characterize the rate-memory tradeoff for a decentralized setting, in which users fill out their cache content without any coordination.

I. Introduction

Caching is a commonly used approach to reduce traffic rate in a network system during peak-traffic times, by duplicating part of the content in the memories distributed across the network. In its basic form, a caching system operates in two phases: (1) a placement phase, where each cache is populated up to its size, and (2) a delivery phase, where the users reveal their requests for content and the server has to deliver the requested content. During the delivery phase, the server exploits the content of the caches to reduce the network traffic.

Conventionally, the caching systems have been based on uncoded unicast delivery where the objective is mainly to maximize the hit rate, i.e. the chance the requested content can be delivered locally [1]–[8]. While in systems with single cache memory this approach can achieve optimal performance, it has been recently shown in [9] that for multi-cache systems, the optimality no longer holds. In [9], an information theoretic framework for multi-cache systems was introduced, and it was shown that coding can offer a significant gain that scales with the size of the network. Several coded caching schemes have been proposed since then [10]–[15]. The caching problem has also been extended in various directions, including decentralized caching [16], online caching [17], caching with nonuniform demands [18]–[20], hierarchical caching [21]–[23], Device-to-device caching [24], cache-aided interference channels [25]–[28], caching on file selection networks [29]–[31], caching on broadcast channels [32]–[35], and caching for channels with delayed feedback with channel state information [36], [37]. The same idea is also useful in the context of distributed computing, in order to take advantage of extra computation to reduce the communication load [38]–[40].

Characterizing the exact rate-memory tradeoff in the above caching scenarios is an active line of research. Besides developing better achievability schemes, there have been efforts in tightening the outer bound of the rate-memory tradeoff [13], [31], [41]–[43]. Nevertheless, in almost all scenarios, there is still a gap between the state-of-the-art communication load and the converse, leaving the exact rate-memory tradeoff as an open problem.

In this paper, we focus on an important class of caching schemes, where the prefetching scheme is required to be uncoded. In fact, almost all caching schemes proposed for the above mentioned problems use uncoded prefetching. As a main advantage, uncoded prefetching allows us to handle asynchronous demands without increasing the communication rates, by dividing files into smaller subfiles [16]. Within this class of caching schemes, we characterize the exact rate-memory tradeoff for both the *average rate* for uniform files popularity and the *peak rate*, in both centralized and decentralized settings, for all possible values of the parameters in the problem.

In particular, we first propose a novel caching strategy for the centralized setting. We demonstrate how to exploit commonality among user demands, and reduce the communication load for the more general case where not all user demands are distinct. The proposed scheme strictly improves the state of the art in terms of the expected communication rate and the peak rate.

In addition, we demonstrate the exact optimality of the proposed scheme through a matching converse. The main idea is to divide the set of all demands into smaller subsets (referred to as types), and derive tight lower bounds for the minimum peak rate and the minimum average rate on each type separately. We show that, when the prefetching is uncoded, the exact

rate-memory tradeoff can be completely characterized using this technique, and the placement phase in the proposed caching scheme universally achieves those minimum rates on all types.

Furthermore, we extend the techniques we developed for the centralized caching problem to characterize the exact rate-memory tradeoff in the decentralized setting. Based on the proposed centralized caching scheme, we develop a new decentralized caching scheme that strictly improves the state of the art [15], [16]. In addition, we formally define the framework of decentralized caching, and prove matching converses given the framework, showing that the proposed scheme is optimal.

The problem of caching with uncoded prefetching has also been considered before in [11], [44]. In particular, by introducing a new lower bounds, it was shown that the scheme of [9] is optimal when considering *peak rate* and *centralized caching*, if the number of files is larger than the number of users. In contrary, in this paper we characterize the rate-memory tradeoff for all possible values of the parameters of the problem, for both peak rate and average rate, and both centralized and decentralized settings. To accomplish this, we further improved the outerbound of [11], [44]. More importantly, we introduce a new achieability scheme, which strictly improves the scheme of [9].

The rest of this paper is organized as follows. Section II formally establishes a centralized caching framework, and defines the main problem studied in this paper. Section III summarizes the main result of this paper for the centralized setting. Section IV describes and demonstrates the optimal centralized caching scheme that achieves the minimum expected rate and the minimum peak rate. Section V proves matching converses that show the optimality of the proposed centralized caching scheme. Section VI extends the techniques we developed for the centralized caching problem to characterize the exact rate-memory tradeoff in the decentralized setting.

II. System Model and Problem Definition

In this section, we formally introduce the system model for centralized caching problem. Then, we define the rate memory trade-off based on the introduced framework, and state the main problem studied in this paper.

A. System Model

We consider a system with one server connected to K users through a shared, error-free link (see Fig. 1). The server has access to a database of N files $W_1, ..., W_N$, each of size F bits.[1] We denote the jth bit in file i by $B_{i,j}$, and we assume that all bits in the database are i.i.d. Bernoulli random variables with $p = 0.5$. Each user has an isolated cache memory of size MF bits, where $M \in [0, N]$. For convenience, we define parameter $r = \frac{KM}{N}$.

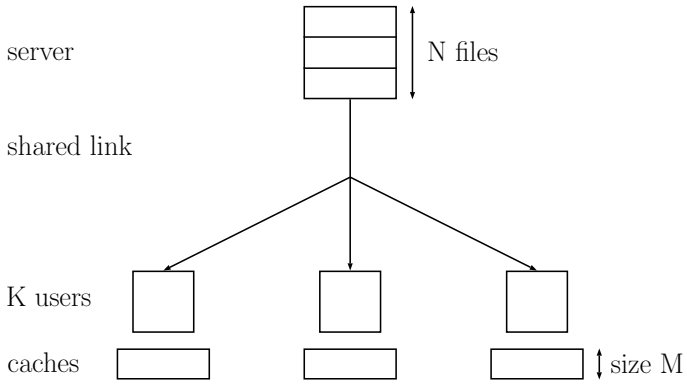

Fig. 1: Caching system considered in this paper. The figure illustrates the case where $K = N = 3$, $M = 1$.

The system operates in two phases: a placement phase and a delivery phase. In the placement phase, users are given access to the entire database, and each user can fill their cache using the database. However, instead of allowing coding in the prefetching [9], we focus on an important class of prefetching schemes, refered to as uncoded prefetching schemes, defined as follows: Each user k selects a subset of size MF bits from the database and stores the values in their own cache. Let $\mathcal{M}_k$ denote the set of selected indices. In the delivery phase, only the server has access to the database of files. Each user k requests one of the files in the database. To characterize the requests from the users, we define *demand* $\boldsymbol{d} = (d_1, ..., d_K)$, where d_k is the index of the file requested by user k. We denote the number of distinct requested files in $\boldsymbol{d}$ by $N_{\mathrm{e}}(\boldsymbol{d})$. We denote the set of all possible demands by $\mathcal{D}$, i.e., $\mathcal{D} = \{1, ..., N\}^K$.

The server is informed of the demand and proceeds by generating a signal X of size RF bits as a function of $W_1, ..., W_N$, and transmits the signal over the shared link. R is a fixed real number given the demand $\boldsymbol{d}$. The quantities RF and R are referred to as the load and the rate of the shared link, respectively. Using the values of bits in $\mathcal{M}_k$ and the signal X received over the shared link, each user k aims to reconstruct its requested file W_{d_k}.

[1] Although we only focus on binary files, the same techniques developed in this paper can also be used for cases of q-ary files and files using a mixture of different alphabets, to prove that same rate-memory trade off holds.

B. Problem Definition

Based on the above framework, we define the rate-memory tradeoff for average rate using the following terminology. Given a prefetching $\boldsymbol{\mathcal{M}} = (\mathcal{M}_1, ..., \mathcal{M}_K)$, we define that a communication rate R is *achievable* for demand $\boldsymbol{d}$ if and only if there exists a message X of length RF such that every active user k is able to recover its desired file W_{d_k}. We denote $R^*(\boldsymbol{d}, \boldsymbol{\mathcal{M}})$ as the minimum achievable rate given $\boldsymbol{d}$ and $\boldsymbol{\mathcal{M}}$.

We assume that all users are making requests independently, and that all files are equally likely to be requested by each user. Thus, the probability distribution of the demand $\boldsymbol{d}$ is uniform on $\mathcal{D}$. We define the average rate $R^*(\boldsymbol{\mathcal{M}})$ as the expected minimum achievable rate given a prefetching $\boldsymbol{\mathcal{M}}$ under uniformly random demand, i.e.,

$$R^*(\boldsymbol{\mathcal{M}}) = \mathbb{E}_{\boldsymbol{d}}[R^*(\boldsymbol{d}, \boldsymbol{\mathcal{M}})]. \tag{1}$$

The rate-memory trade-off for average rate is essentially finding the minimum average rate R^* for different memory constraints M that can be achieved by prefetchings that satisfy this constraint. In other words, we want to find

$$R^* = \min_{\boldsymbol{\mathcal{M}}} R^*(\boldsymbol{\mathcal{M}}). \tag{2}$$

as a function of N, K, M, and for sufficiently large F.

Similarly, the rate-memory tradeoff for peak rate is essentially finding the minimum peak rate, denoted by R^*_{worst}, which is formally defined in appendix B.

III. Main Results

We state the main result of this paper in the following theorem.

Theorem 1. *For a caching problem with K users, a database of N files, local cache size of M files at each user, and parameter $r = \frac{KM}{N}$, we have*

$$R^* = \mathbb{E}_{\boldsymbol{d}}\left[\frac{\binom{K}{r+1} - \binom{K-N_{\text{e}}(\boldsymbol{d})}{r+1}}{\binom{K}{r}}\right], \tag{3}$$

for $r \in \{0, 1, ..., K\}$, where $\boldsymbol{d}$ is uniformly random in $\mathcal{D} = \{1, ..., N\}^K$ and $N_{\text{e}}(\boldsymbol{d})$ denotes the number of distinct requests in $\boldsymbol{d}$. Furthermore, for $r \notin \{0, 1, ..., K\}$, R^* equals the lower convex envelope of its values at $r \in \{0, 1, ..., K\}$.[2]

Remark 1. To prove Theorem 1, we propose a new caching scheme that strictly improves the state of the art [9], which was relied on by all prior works considering the minimum average rate for the caching problem [18], [30], [43]. In particular, the rate achieved by the previous best known caching scheme equals the lower convex envelope of $\min\{\frac{K-r}{r+1}, \mathbb{E}_{\boldsymbol{d}}[N_{\text{e}}(\boldsymbol{d})(1-\frac{r}{K})]\}$ at $r \in \{0, 1, ..., K\}$, which is strictly larger than R^* when $N > 1$ and $r < K - 1$. For example, when $K = 30$, $N = 30$, and $r = 1$, the state-of-the-art scheme requires a communication rate of 14.12, while the proposed scheme achieves the rate 12.67.

The improvement of our proposed scheme over the state of the art can be interpreted intuitively as follows. The caching scheme proposed in [9] essentially decomposes the problem into 2 scenarios: in one case, the redundancy of user demands is ignored, and the information is delivered by satisfying different demands using single coded multicast transmission; in the other case, random coding is used to deliver the same request to multiple receivers. Our result demonstrates that the decomposition of the caching problem into these 2 scenarios is suboptimal, and our proposed caching scheme precisely accounts for the effect of redundant user demands.

Remark 2. The technique for finding the minimum average rate in the centralized setting can be straightforwardly extended to find the minimum peak rate, which was solved for $N \geq K$ [44]. Here we show that we not only recover their result, but also fully characterize the rate for all possible values of N and K, resulting in the following corollary, which will be proved in Appendix B.

Corollary 1. *For a caching problem with K users, a database of N files, a local cache size of M files at each user, and parameter $r = \frac{KM}{N}$, we have*

$$R^*_{\text{worst}} = \frac{\binom{K}{r+1} - \binom{K-\min\{K,N\}}{r+1}}{\binom{K}{r}} \tag{4}$$

for $r \in \{0, 1, ..., K\}$. Furthermore, for $r \notin \{0, 1, ..., K\}$, R^*_{worst} equals the lower convex envelope of its values at $r \in \{0, 1, ..., K\}$.

Remark 3. As we will discuss in Section VI, we can also extend the techniques that we developed for proving Theorem 1 to the decentralized setting. The exact rate-memory tradeoff for both the average rate and the peak rate can be fully characterized using these techniques. Besides, the newly proposed decentralized caching scheme for achieving the minimum rates strictly improves the state of the art [15], [16].

In the following sections, we prove Theorem 1 by first describing a caching scheme that achieves the minimum average rate (see Section IV), and then deriving tight lower bounds of the expected rates for any uncoded prefetching scheme (see Section V).

[2]In this paper we define $\binom{n}{k} = 0$ when $k > n$.

IV. The Optimal Caching Scheme

In this section, we provide a caching scheme (i.e. a prefetching scheme and a delivery scheme) to achieve R^* stated in Theorem 1. Before introducing the proposed caching scheme, we demonstrate the main ideas of the proposed scheme through a motivating example.

A. Motivating Example

Consider a caching system with 3 files (denoted by A, B, and C), 6 users, and a caching size of 1 file for each user. To develop a caching scheme, we need to design an uncoded prefetching scheme, independent of the demands, and develop delivery strategies for each of the possible 3^6 demands.

For the prefetching strategy, we break file A into 15 subfiles of equal size, and denote their values by $A_{\{1,2\}}$, $A_{\{1,3\}}$, $A_{\{1,4\}}$, $A_{\{1,5\}}$, $A_{\{1,6\}}$, $A_{\{2,3\}}$, $A_{\{2,4\}}$, $A_{\{2,5\}}$, $A_{\{2,6\}}$, $A_{\{3,4\}}$, $A_{\{3,5\}}$, $A_{\{3,6\}}$, $A_{\{4,5\}}$, $A_{\{4,6\}}$, and $A_{\{5,6\}}$. Each user k caches the subfiles whose second index includes k, e.g., user 1 caches $A_{\{1,2\}}$, $A_{\{1,3\}}$, $A_{\{1,4\}}$, $A_{\{1,5\}}$, and $A_{\{1,6\}}$. The same goes for files B and C. This prefetching scheme was originally proposed in [9].

Given the above prefetching scheme, we now need to develop an optimal delivery strategy for each of the possible demands. In this subsection, we demonstrate the key idea of our proposed delivery scheme through a representative demand scenario, namely, each file is requested by 2 users as shown in Figure 2.

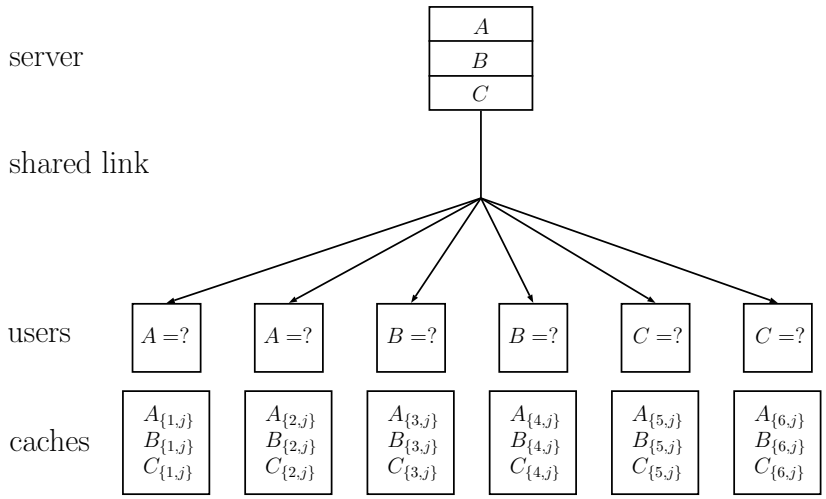

Fig. 2: A caching system with 6 users, 3 files, local cache size pf 1 file at each user, and a demand where each file is requested by 2 users.

We first consider a subset of 3 users $\{1,2,3\}$. User 1 requires subfile $A_{\{2,3\}}$, which is only available at users 2 and 3. User 2 requires subfile $A_{\{1,3\}}$, which is only available at users 1 and 3. User 3 requires subfile $B_{\{1,2\}}$, which is only available at users 1 and 2. In other words, the three users would like to exchange subfiles $A_{\{2,3\}}$, $A_{\{1,3\}}$, and $B_{\{1,2\}}$, which can be enabled by transmitting the message $A_{\{2,3\}} \oplus A_{\{1,3\}} \oplus B_{\{1,2\}}$ over the shared link.

Similarly, we can create and broadcast messages for any subset $\mathcal{A}$ of 3 users that exchange 3 subfiles among those 3 users. As a short hand notation, we denote the corresponding message by $Y_{\mathcal{A}}$. According to the delivery scheme proposed in [9], if we broadcast all $\binom{6}{3} = 20$ messages that could be created in this way, all users are able to decode their requested files.

However, in this paper we propose a delivery scheme that, instead of broadcasting all those 20 messages, only 19 of them are computed and broadcasted, omitting the message $Y_{\{2,4,6\}}$. Specifically, we broadcast the following 19 values:

$$Y_{\{1,2,3\}} = B_{\{1,2\}} \oplus A_{\{1,3\}} \oplus A_{\{2,3\}} \qquad Y_{\{1,2,4\}} = B_{\{1,2\}} \oplus A_{\{1,4\}} \oplus A_{\{2,4\}}$$
$$Y_{\{1,2,5\}} = C_{\{1,2\}} \oplus A_{\{1,5\}} \oplus A_{\{2,5\}} \qquad Y_{\{1,2,6\}} = C_{\{1,2\}} \oplus A_{\{1,6\}} \oplus A_{\{2,6\}}$$
$$Y_{\{1,3,4\}} = B_{\{1,3\}} \oplus B_{\{1,4\}} \oplus A_{\{3,4\}} \qquad Y_{\{1,3,5\}} = C_{\{1,3\}} \oplus B_{\{1,5\}} \oplus A_{\{3,5\}}$$
$$Y_{\{1,3,6\}} = C_{\{1,3\}} \oplus B_{\{1,6\}} \oplus A_{\{3,6\}} \qquad Y_{\{1,4,5\}} = C_{\{1,4\}} \oplus B_{\{1,5\}} \oplus A_{\{4,5\}}$$
$$Y_{\{1,4,6\}} = C_{\{1,4\}} \oplus B_{\{1,6\}} \oplus A_{\{4,6\}} \qquad Y_{\{1,5,6\}} = C_{\{1,5\}} \oplus C_{\{1,6\}} \oplus A_{\{5,6\}}$$
$$Y_{\{2,3,4\}} = B_{\{2,3\}} \oplus B_{\{2,4\}} \oplus A_{\{3,4\}} \qquad Y_{\{2,3,5\}} = C_{\{2,3\}} \oplus B_{\{2,5\}} \oplus A_{\{3,5\}}$$
$$Y_{\{2,3,6\}} = C_{\{2,3\}} \oplus B_{\{2,6\}} \oplus A_{\{3,6\}} \qquad Y_{\{2,4,5\}} = C_{\{2,4\}} \oplus B_{\{2,5\}} \oplus A_{\{4,5\}}$$
$$Y_{\{2,5,6\}} = C_{\{2,5\}} \oplus C_{\{2,6\}} \oplus A_{\{5,6\}} \qquad Y_{\{3,4,5\}} = C_{\{3,4\}} \oplus B_{\{3,5\}} \oplus B_{\{4,5\}}$$
$$Y_{\{3,4,6\}} = C_{\{3,4\}} \oplus B_{\{3,6\}} \oplus B_{\{4,6\}} \qquad Y_{\{3,5,6\}} = C_{\{3,5\}} \oplus C_{\{3,6\}} \oplus B_{\{5,6\}}$$
$$Y_{\{4,5,6\}} = C_{\{4,5\}} \oplus C_{\{4,6\}} \oplus B_{\{5,6\}}$$

Surprisingly, even after taking out the extra message, all users are still able to decode the requested files. The reason is as follows:

User 1 is able to decode file A, because every subfile $A_{\{i,j\}}$ that is not cached by user 1 can be computed with the help of $Y_{\{1,i,j\}}$, which is directly broadcasted. The above is the same decoding procedure used in [9]. User 2 can easily decode all subfiles in A except $A_{\{4,6\}}$ in a similar way, although decoding $A_{\{4,6\}}$ is more challenging since the value $Y_{\{2,4,6\}}$, which is needed in the above decoding procedure for decoding $A_{\{4,6\}}$, is not directly broadcasted. However, user 2 can still decode $A_{\{4,6\}}$ by adding $Y_{\{1,4,6\}}$, $Y_{\{1,4,5\}}$, $Y_{\{1,3,6\}}$, and $Y_{\{1,3,5\}}$, which gives the binary sum of $A_{\{4,6\}}$, $A_{\{4,5\}}$, $A_{\{3,6\}}$, and $A_{\{3,5\}}$. Because $A_{\{4,5\}}$, $A_{\{3,6\}}$, and $A_{\{3,5\}}$ are easily decodable, $A_{\{4,6\}}$ can be obtained consequently.

Due to symmetry, all other users can decode their requested files in the same manner. This completes the decoding tasks for the given demand.

B. General Schemes

Now we present a general caching scheme that achieves the rate R^* stated in Theorem 1. We focus on presenting prefetching schemes and delivery schemes when $r \in \{0, 1, ..., K\}$, since for general r, the minimum rate R^* can be achieved by memory sharing.

Remark 4. Note that the rates stated in equation (3) for $r \in \{0, 1, ..., K\}$ form a convex sequence, which are consequently on their lower convex envelope. Thus those rates cannot be further improved using memory sharing.

To prove the achievability of R^*, we consider the following optimal prefetching: We partition each file i into $\binom{K}{r}$ non-overlapping subfiles with approximately equal size. We assign the $\binom{K}{r}$ subfiles to $\binom{K}{r}$ different subsets of $\{1, .., K\}$ of size r, and we denote the value of the subfile assigned to subset $\mathcal{A}$ by $W_{i,\mathcal{A}}$. Given this partition, each user k cache all bits in all subfiles $W_{i,\mathcal{A}}$ that satisfy $k \in \mathcal{A}$. Because each user caches $\binom{K-1}{r-1}N$ subfiles, and each subfile has $F/\binom{K}{r}$ bits, the caching load of each user equals $NrF/K = MF$ bits, which satisfies the memory constraint. This prefetching was originally proposed in [9]. In the rest of the paper, we refer to this prefetching as *symmetric batch prefetching*.

Given this prefetching (denoted by $\boldsymbol{\mathcal{M}}$), our goal is to show that for any demand $\boldsymbol{d}$, we can find a delivery scheme that achieves the following optimal rate:

$$R^*(\boldsymbol{d}, \boldsymbol{\mathcal{M}}) = \frac{\binom{K}{r+1} - \binom{K-N_{\mathrm{e}}(\boldsymbol{d})}{r+1}}{\binom{K}{r}}. \tag{5}$$

Hence, by taking the expectation about the demand $\boldsymbol{d}$, the rate R^* stated in Theorem 1 can be achieved.

Remark 5. Note that, in the special case where all users are requesting different files (i.e., $N_{\mathrm{e}}(\boldsymbol{d}) = K$), the above rate equals $\frac{K-r}{r+1}$, which can already be achieved by the delivery scheme proposed in [9]. Our proposed scheme aims to achieve this optimal rate in more general circumstances, when some users may share common demands.

Remark 6. Finding the minimum communication load given a prefetching $\boldsymbol{\mathcal{M}}$ can be viewed as a special case of the index coding problem. Theorem 1 indicates the optimality of the delivery scheme given the symmetric batch prefetching, which implies that (5) gives the solution to a special class of non-symmetric index coding problem.

The optimal delivery scheme is designed as follows: For each demand $\boldsymbol{d}$, recall that $N_{\mathrm{e}}(\boldsymbol{d})$ denotes the number of distinct files requested by all users. The server arbitrarily selects a subset of $N_{\mathrm{e}}(\boldsymbol{d})$ users, denoted by $\mathcal{U} = \{u_1, ..., u_{N_{\mathrm{e}}(\boldsymbol{d})}\}$, that requests $N_{\mathrm{e}}(\boldsymbol{d})$ different files. We refer to these users as *leaders*.

Given an arbitrary subset $\mathcal{A}$ of $r+1$ users, each user $k \in \mathcal{A}$ needs the subfile $W_{d_k,\mathcal{A}\backslash\{k\}}$, which is known by all other users in $\mathcal{A}$. In other words, all users in set $\mathcal{A}$ would like to exchange subfiles $W_{d_k,\mathcal{A}\backslash\{k\}}$ for all $k \in \mathcal{A}$. This subfile exchanging can be processed if the binary sum of all those files, i.e. $\underset{x\in\mathcal{A}}{\oplus} W_{d_x,\mathcal{A}\backslash\{x\}}$, is available from the broadcasted message. To simplify the description of the delivery scheme, for each subset $\mathcal{A}$ of users, we define the following short hand notation

$$Y_{\mathcal{A}} = \underset{x\in\mathcal{A}}{\oplus} W_{d_x,\mathcal{A}\backslash\{x\}}. \tag{6}$$

To achieve the rate stated in (5), the server only greedily broadcasts the binary sums that directly helps at least 1 leader. Rigorously, the server computes and broadcasts all $Y_{\mathcal{A}}$ for all subset $\mathcal{A}$ of size $r+1$ that satisfies $\mathcal{A} \cap \mathcal{U} \neq \emptyset$. The length of the message equals $\binom{K}{r+1} - \binom{K-N_{\mathrm{e}}(\boldsymbol{d})}{r+1}$ times the size of a subfile, which matches the stated rate.

We now argue that each user who requests a file is able to decode the requested file upon receiving the messages. For any leader $k \in \mathcal{U}$ and any subfile $W_{d_k,\mathcal{A}}$ that is requested but not cached by user k, the message $Y_{\{k\}\cup\mathcal{A}}$ is directly available from the broadcast. Thus, k is able to obtain all requested subfiles by decoding each subfile $W_{d_k,\mathcal{A}}$ from message $Y_{\{k\}\cup\mathcal{A}}$.

The decoding procedure for a non-leader user k is less straight forward, because not all messages $Y_{\{k\}\cup\mathcal{A}}$ for corresponding required subfiles $W_{d_k,\mathcal{A}}$ are directly broadcasted. However, user k can generate these messages simply based on the received messages, and can thus decode all required subfiles. We prove the above fact as follows:

First we prove the following simple lemma:

Lemma 1. *Given a demand $\boldsymbol{d}$, and a set of leaders $\mathcal{U}$. For any subset $\mathcal{B} \subseteq \{1, ..., K\}$ that includes $\mathcal{U}$, let $\mathcal{V}_{\mathrm{F}}$ be family of all subsets $\mathcal{V}$ of $\mathcal{B}$ such that each requested file in $\boldsymbol{d}$ is requested by exactly one user in $\mathcal{V}$.*

The following equation holds:

$$\bigoplus_{\mathcal{V}\in\mathcal{V}_{\mathrm{F}}} Y_{\mathcal{B}\backslash\mathcal{V}} = 0 \tag{7}$$

where $Y_{\mathcal{A}}$ is defined in (6).

Proof. For each $u \in \mathcal{U}$ we define $\mathcal{B}_u$ as follows,

$$\mathcal{B}_u = \{x \in \mathcal{B} \mid d_x = d_u\}. \tag{8}$$

Then all sets $\mathcal{B}_u$ disjointly cover the set $\mathcal{B}$, and the following equations hold:

$$\bigoplus_{\mathcal{V}\in\mathcal{V}_{\mathrm{F}}} Y_{\mathcal{B}\backslash\mathcal{V}} = \bigoplus_{\mathcal{V}\in\mathcal{V}_{\mathrm{F}}} \bigoplus_{x\in\mathcal{B}\backslash\mathcal{V}} W_{d_x,\mathcal{B}\backslash(\mathcal{V}\cup\{x\})} \tag{9}$$

$$= \bigoplus_{u\in\mathcal{U}} \bigoplus_{\mathcal{V}\in\mathcal{V}_{\mathrm{F}}} \bigoplus_{x\in(\mathcal{B}\backslash\mathcal{V})\cap\mathcal{B}_u} W_{d_u,\mathcal{B}\backslash(\mathcal{V}\cup\{x\})} \tag{10}$$

$$= \bigoplus_{u\in\mathcal{U}} \bigoplus_{\mathcal{V}\in\mathcal{V}_{\mathrm{F}}} \bigoplus_{x\in\mathcal{B}_u\backslash\mathcal{V}} W_{d_u,\mathcal{B}\backslash(\mathcal{V}\cup\{x\})}. \tag{11}$$

For each $u \in \mathcal{U}$, we let $\mathcal{V}_u$ be the family of all subsets $\mathcal{V}$ of $\mathcal{B}\backslash\mathcal{B}_u$ such that each requested file in $\boldsymbol{d}$, except d_u, is requested by exactly one user in $\mathcal{V}$. Then $\mathcal{V}_{\mathrm{F}}$ can be represented as follows:

$$\mathcal{V}_{\mathrm{F}} = \{\{y\} \cup \mathcal{V} \mid y \in \mathcal{B}_u, \mathcal{V} \in \mathcal{V}_u\}. \tag{12}$$

Consequently, the following equation holds for each $u \in \mathcal{U}$:

$$\bigoplus_{\mathcal{V}\in\mathcal{V}_{\mathrm{F}}} \bigoplus_{x\in\mathcal{B}_u\backslash\mathcal{V}} W_{d_u,\mathcal{B}\backslash(\mathcal{V}\cup\{x\})} = \bigoplus_{\mathcal{V}\in\mathcal{V}_u} \bigoplus_{y\in\mathcal{B}_u} \bigoplus_{x\in\mathcal{B}_u\backslash\{y\}} W_{d_u,\mathcal{B}\backslash(\mathcal{V}\cup\{x,y\})} \tag{13}$$

$$= 0 \tag{14}$$

Thus the LHS of equation (7) equals 0. □

Consider any subset $\mathcal{A}$ of $r+1$ non-leader users. From Lemma 1, the message $Y_{\mathcal{A}}$ can be directly computed from the broadcasted messages using the following equation:

$$Y_{\mathcal{A}} = \bigoplus_{\mathcal{V}\in\mathcal{V}_{\mathrm{F}}\backslash\{\mathcal{U}\}} Y_{\mathcal{B}\backslash\mathcal{V}}, \tag{15}$$

where $\mathcal{B} = \mathcal{A} \cup \mathcal{U}$, given the fact that all messages on the RHS of the above equation are directly broadcasted. Hence, each user k can obtain the value $Y_{\mathcal{A}}$ for any subset $\mathcal{A}$ of $r+1$ users, and can subsequently decode its requested file as previously discussed.

V. Converse

In this section, we derive a tight lower bound on the minimum expected rate R^*, which shows the optimality of the caching scheme proposed in this paper. To derive the corresponding lower bound on the average rate over all demands, we divide the set $\mathcal{D}$ into smaller subsets, and lower bound the average rates within each subset individually. We refer to these smaller subsets as *types*, which are defined as follows.[3]

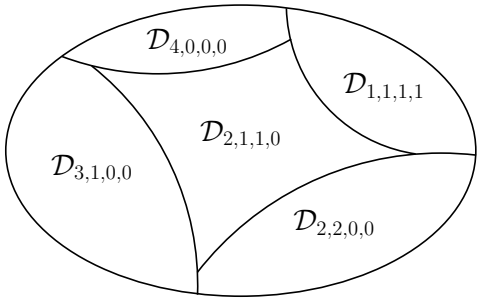

Fig. 3: Dividing $\mathcal{D}$ into 5 types, for a caching problem with 4 files and 4 users.

Given an arbitrary demand $\boldsymbol{d}$, we define its *statistics*, denoted by $\boldsymbol{s}(\boldsymbol{d})$, as a sorted array of length N, such that $s_i(\boldsymbol{d})$ equals the number of users that request the ith most requested file. We denote the set of all possible statistics by $\mathcal{S}$. Grouping by the same statistics, the set of all demands $\mathcal{D}$ can be broken into many small subsets. For any statistics $\boldsymbol{s} \in \mathcal{S}$, we define type $\mathcal{D}_{\boldsymbol{s}}$ as the set of queries with statistics $\boldsymbol{s}$.

For example, consider a caching problem with 4 files (denoted by A, B, C, and D) and 4 users. The statistics of the demand $\boldsymbol{d} = (A, A, B, C)$ equals $\boldsymbol{s}(\boldsymbol{d}) = (2, 1, 1, 0)$. More generally, the set of all possible statistics for this problem is

[3]The notion of type was also recently introduced in [45] in order to simplify the LP for finding better converse bounds for the coded caching problem.

$\mathcal{S} = \{(4,0,0,0),(3,1,0,0),(2,2,0,0),(2,1,1,0),(1,1,1,1)\}$, and $\mathcal{D}$ can be divided into 5 types accordingly, as shown in Fig. 3.

Note that for each demand $\boldsymbol{d}$, the value $N_{\mathrm{e}}(\boldsymbol{d})$ only depends on its statistics $\boldsymbol{s}(\boldsymbol{d})$, and thus the value is identical across all demands in $\mathcal{D}_{\boldsymbol{s}}$. For convenience, we denote that value by $N_{\mathrm{e}}(\boldsymbol{s})$.

Given a prefetching $\boldsymbol{\mathcal{M}}$, we denote the average rate within each type $\mathcal{D}_{\boldsymbol{s}}$ by $R^*(\boldsymbol{s},\boldsymbol{\mathcal{M}})$. Rigorously,

$$R^*(\boldsymbol{s},\boldsymbol{\mathcal{M}}) = \frac{1}{|\mathcal{D}_{\boldsymbol{s}}|}\sum_{\boldsymbol{d}\in\mathcal{D}_{\boldsymbol{s}}} R^*(\boldsymbol{d},\boldsymbol{\mathcal{M}}). \tag{16}$$

Recall that all demands are equally likely, we have

$$R^* = \min_{\boldsymbol{\mathcal{M}}} \mathbb{E}_{\boldsymbol{s}}[R^*(\boldsymbol{s},\boldsymbol{\mathcal{M}})] \tag{17}$$

$$\geq \mathbb{E}_{\boldsymbol{s}}[\min_{\boldsymbol{\mathcal{M}}} R^*(\boldsymbol{s},\boldsymbol{\mathcal{M}})]. \tag{18}$$

Hence, in order to lower bound R^*, it is sufficient to bound the minimum value of $R^*(\boldsymbol{s},\boldsymbol{\mathcal{M}})$ for each type $\mathcal{D}_{\boldsymbol{s}}$ individually. We show that, when the prefetching is uncoded, the minimum rate within a type can be tightly bounded, thus the exact rate-memory tradeoff can be completely characterized using this technique.

The tight lower bounds of the minimum average rates within each type are presented in the following lemma:

Lemma 2. *Consider a caching problem with N files, K users, and a local cache size of M files for each user. For any type $\mathcal{D}_{\boldsymbol{s}}$, the minimum value of $R^*(\boldsymbol{s},\boldsymbol{\mathcal{M}})$ is lower bounded by*

$$\min_{\boldsymbol{\mathcal{M}}} R^*(\boldsymbol{s},\boldsymbol{\mathcal{M}}) \geq \mathrm{Conv}\left(\frac{\binom{K}{r+1} - \binom{K-N_{\mathrm{e}}(\boldsymbol{s})}{r+1}}{\binom{K}{r}}\right) \tag{19}$$

where $\mathrm{Conv}(f(r))$ *denotes the lower convex envelope of the following points:* $\{(r,f(r)) \mid r\in\{0,1,...,K\}\}$.

Remark 7 (Universal Optimality of Symmetric Batch Prefetching). The above lemma characterizes the minimum average rate given a type $\mathcal{D}_{\boldsymbol{s}}$, if the prefetching $\boldsymbol{\mathcal{M}}$ can be designed based on s. However, for (18) to be tight, the average rate for each different type has to be minimized on the same prefetching. Surprisingly, such an optimal prefetching exists, an example being the symmetric batch prefetching according to Section IV. This indicates that the symmetric batch prefetching is universally optimal for all types in terms of the average rates.

Remark 8. Lemma 2 can be viewed as an extension of the converses in [11], [44] to the scenario where a more general subset of demands are considered: only a proper subset of files are requested, and not all user demands are distinct.

We postpone the proof of Lemma 2 to Appendix A and first prove the converse using the lemma.

From (18) and Lemma 2, R^* can be lower bounded as follows:

$$R^* \geq \mathbb{E}_{\boldsymbol{s}}[\min_{\boldsymbol{\mathcal{M}}} R^*(\boldsymbol{s},\boldsymbol{\mathcal{M}})] \tag{20}$$

$$\geq \mathbb{E}_{\boldsymbol{s}}\left[\mathrm{Conv}\left(\frac{\binom{K}{r+1} - \binom{K-N_{\mathrm{e}}(\boldsymbol{s})}{r+1}}{\binom{K}{r}}\right)\right]. \tag{21}$$

Because the sequence

$$c_n = \frac{\binom{K}{n+1} - \binom{K-N_{\mathrm{e}}(\boldsymbol{s})}{n+1}}{\binom{K}{n}} \tag{22}$$

is convex, we can switch the order of the expectation and the Conv in (21). Therefore, R^* is lower bounded by the rate defined in Theorem 1.[4]

VI. Extension to the Decentralized Setting

In the sections above, we introduced a new centralized caching scheme and a new bounding technique that completely characterize the minimum average communication rate and the minimum peak rate, when the prefetching is required to be uncoded. Interestingly, these techniques can also be extended to fully characterize the rate-memory tradeoff for decentralized caching. In this section, we formally establish a system model for decentralized caching systems, and state the exact rate-memory tradeoff as main results for both the average rate and the peak rate.

A. System Model and Problem Formulation

In many practical systems, out of the large number of users that may potentially request files from the server through the shared error-free link, only a random unknown subset are connected to the link and making requests at any given time instance. To handle this situation, the concept of decentralized prefetching scheme was introduced in [16], where each user has to fill their caches randomly and independently, based on the same probability distribution.

[4]As noted in remark 4, the rate R^* stated in equation (3) for $r\in\{0,1,...,K\}$ is convex, so it is sufficient to prove R^* is lower bounded by the convex envelope of its values at $r\in\{0,1,...,K\}$.

Based on the above framework, we formally define decentralized caching as follows: Instead of following a deterministic caching scheme, each user k caches a subset $\mathcal{M}_k$ of size MF bits randomly and independently. A *decentralized prefetching scheme*, denoted by $P_{\boldsymbol{\mathcal{M}}}$, is defined as a sequence of probability distribution of the prefetching $\boldsymbol{\mathcal{M}}$, parameterized by F.

When K users are making requests, we say that a rate R is *achievable* given a prefetching scheme $P_{\boldsymbol{\mathcal{M}}}$ and a demand $\boldsymbol{d} \in \{1, ..., N\}^K$, if for any $\epsilon > 0$, the rate $R^*(\boldsymbol{d}, \boldsymbol{\mathcal{M}})$ is no greater than $R + \epsilon$ with high probability for large F. We denote the minimum of such rates by $R^*_K(\boldsymbol{d}, P_{\boldsymbol{\mathcal{M}}})$.

For the average rate, we can define the rate-memory tradeoff as follows. Given the fact that a decentralized prefetching scheme is designed without the knowledge of the number of active users K, we characterize the performance of a prefetching scheme $P_{\boldsymbol{\mathcal{M}}}$ using an infinite dimensional vector $R^*_K(P_{\boldsymbol{\mathcal{M}}})$, which is defined as

$$R^*_K(P_{\boldsymbol{\mathcal{M}}}) = \mathbb{E}_{\boldsymbol{d}}[R^*_K(\boldsymbol{d}, P_{\boldsymbol{\mathcal{M}}})], \qquad \forall K \in \mathbb{N}, \tag{23}$$

where the demand $\boldsymbol{d}$ given each K is uniformly distributed on $\{1, ..., N\}^K$.

We aim to find the region in the infinite dimensional vector space that can be achieved by any decentralized prefetching scheme under the memory constraint, denoted by $\mathcal{R}$, and defined as

$$\mathcal{R} = \bigcup_{P_{\boldsymbol{\mathcal{M}}}} \{\{R_K\}_{K\in\mathbb{N}} \mid \forall K \in \mathbb{N}, R_K \geq R^*_K(P_{\boldsymbol{\mathcal{M}}})\}, \tag{24}$$

as a function of N and M.

Similarly, we characterize the peak rate performance of a prefetching scheme $P_{\boldsymbol{\mathcal{M}}}$ using an infinite dimensional vector $R^*_{K,\text{worst}}(P_{\boldsymbol{\mathcal{M}}})$, which is defined as

$$R^*_{K,\text{worst}}(P_{\boldsymbol{\mathcal{M}}}) = \max_{\boldsymbol{d}\in\mathcal{D}} R^*_K(\boldsymbol{d}, P_{\boldsymbol{\mathcal{M}}}), \qquad \forall K \in \mathbb{N}. \tag{25}$$

We aim to find the region in the infinite dimensional vector space that can be achieved by any decentralized prefetching scheme under the memory constraint, denoted by $\mathcal{R}_{\text{worst}}$, and defined as

$$\mathcal{R}_{\text{worst}} = \bigcup_{P_{\boldsymbol{\mathcal{M}}}} \{\{R_K\}_{K\in\mathbb{N}} \mid \forall K \in \mathbb{N}, R_K \geq R^*_{K,\text{worst}}(P_{\boldsymbol{\mathcal{M}}})\}, \tag{26}$$

as a function of N and M.

B. Exact Rate-Memory Tradeoff for Decentralized Setting

The following theorem completely characterizes the exact rate-memory tradeoff for the average rate in decentralized setting:

Theorem 2. *For a decentralized caching problem with parameters N and M, $\mathcal{R}$ is completely characterized by the following equation:*

$$\mathcal{R} = \{\{R_K\}_{K\in\mathbb{N}} \mid R_K \geq \mathbb{E}_{\boldsymbol{d}}[\frac{N-M}{M}(1 - (\frac{N-M}{N})^{N_{\text{e}}(\boldsymbol{d})})]\}, \tag{27}$$

where demand $\boldsymbol{d}$ given each K is uniformly distributed on $\{1, ..., N\}^K$ and $N_{\text{e}}(\boldsymbol{d})$ denotes the number of distinct requests in $\boldsymbol{d}$.[5]

The proof of the above theorem is provided in Appendix C.

Remark 9. Theorem 2 demonstrates that $\mathcal{R}$ has a very simple shape with one dominating point: $\{R_K = \mathbb{E}_{\boldsymbol{d}}[\frac{N-M}{M}(1 - (\frac{N-M}{N})^{N_{\text{e}}(\boldsymbol{d})})]\}_{K\in\mathbb{N}}$. In other words, we can find a decentralized prefetching scheme that simultaneously achieves the minimum expected rates for all possible numbers of active users. Therefore, there is no tension among the expected rates for different numbers of active users. In Appendix C, we will show that one example of the optimal prefetching scheme is to let each user cache $\frac{MF}{N}$ bits in each file uniformly independently.

Remark 10. To prove Theorem 2, we propose a decentralized caching scheme that strictly improves the state of the art [15], [16] (see Appendix C-A), for both the average rate and the peak rate. In particular for the average rate, the state-of-the-art-scheme proposed in [16] achieves the rate $\frac{N-M}{N} \cdot \min\{\frac{N}{M}(1 - (1 - \frac{M}{N})^K), \mathbb{E}_{\boldsymbol{d}}[N_e(\boldsymbol{d})]\}$, which is strictly larger than the rate achieved by the proposed scheme $\mathbb{E}_{\boldsymbol{d}}[\frac{N-M}{M}(1 - (\frac{N-M}{N})^{N_{\text{e}}(\boldsymbol{d})})]$ in most cases.

Remark 11. We also prove a matching information-theoretic outer bound of $\mathcal{R}$ by showing that, the achievable rate of any decentralized caching scheme can be lower bounded by the achievable rate of a caching scheme with centralized prefetching that is used on a system where, there are a large number of users that may potentially request a file, but only a subset of K users are actually making the request. Interestingly, the tightness of this bound indicates that, in a system where the number of potential users is significantly larger than the number of active users, focusing on decentralized prefetching schemes, i.e., letting each user cache bits according to an i.i.d., is optimal.

Using the proposed decentralized caching scheme and the same converse bounding technique, the following corollary, which completely characterizes the exact rate-memory tradeoff for the peak rate in decentralized setting, directly follows:

[5]if $M = 0$, $\mathcal{R} = \{\{R_K\}_{K\in\mathbb{N}} \mid R_K \geq \mathbb{E}_{\boldsymbol{d}}[N_{\text{e}}(\boldsymbol{d})]\}$.

Corollary 2. *For a decentralized caching problem with parameters N and M, the achievable region $\mathcal{R}_{\text{worst}}$ is completely characterized by the following equation:*[6]

$$\mathcal{R}_{\text{worst}} = \{\{R_K\}_{K\in\mathbb{N}} \mid R_K \geq \frac{N-M}{M}(1-(\frac{N-M}{N})^{\min\{N,K\}})\}. \tag{28}$$

The proof of the above corollary is provided in Appendix D.

Remark 12. Corollary 2 demonstrates that $\mathcal{R}_{\text{worst}}$ has a very simple shape with one dominating point: $\{R_K = \frac{N-M}{M}(1-(\frac{N-M}{N})^{\min\{N,K\}})\}_{K\in\mathbb{N}}$. In other words, we can find a decentralized prefetching scheme that simultaneously achieves the minimum peak rates for all possible numbers of active users. Therefore, there is no tension among the peak rates for different numbers of active users. In Appendix D, we will show that one example of the optimal prefetching scheme is to let each user cache $\frac{MF}{N}$ bits in each file uniformly independently.

Remark 13. Similar to the average rate case, a matching converse can be proved by deriving the minimum achievable rates of centralized caching schemes in a system where only a subset of users are actually making the request. Consequently, in a caching system where the number of potential users is significantly larger than the number of active users, focusing on decentralized prefetching schemes is also optimal in terms of the peak rate.

Appendix A
Proof of Lemma 2

Proof. We first use a genie-aided approach to derive a lower bound of $R^*(\boldsymbol{d},\boldsymbol{\mathcal{M}})$ for any demand $\boldsymbol{d}$ and for any prefetching $\boldsymbol{\mathcal{M}}$: Given a demand $\boldsymbol{d}$, let $\mathcal{U}=\{u_1,...,u_{N_{\text{e}}(\boldsymbol{d})}\}$ be an arbitrary subset of $N_{\text{e}}(\boldsymbol{d})$ users that request distinct files. We construct a virtual user whose cache is initially empty. Suppose for each $\ell\in\{1,...,N_{\text{e}}(\boldsymbol{d})\}$, a genie fills the cache with the value of bits that are cached by u_ℓ, but not from files requested by users in $\{u_1,...,u_{\ell-1}\}$. Then with all the cached information provided by the genie, the virtual user should be able to inductively decode all files requested by users in $\mathcal{U}$ upon receiving the message X. Hence, the minimum length of the message is lower bounded by the number of bits in the decodable files, that are not provided by the genie:

$$R^*(\boldsymbol{d},\mathcal{M}_{:})F \geq \sum_{i=1}^{N_{\text{e}}(\boldsymbol{d})}\sum_{j=1}^{F}\mathbb{1}(B_{d_{u_i},j} \text{ is not cached by any user in } \{u_1,...,u_i\}) \tag{29}$$

$$\geq \sum_{i=1}^{N_{\text{e}}(\boldsymbol{d})}\sum_{j=1}^{F}\mathbb{1}((d_{u_i},j)\notin\bigcup_{\ell=1}^{i}\mathcal{M}_{u_\ell}). \tag{30}$$

Using the above bound, we derive a lower bound of the average rates as follows: For any positive integer i, we denote the set of all permutations on $\{1,...,i\}$ by $\mathcal{P}_i$. Then, for each $p_1\in\mathcal{P}_K$ and $p_2\in\mathcal{P}_N$ given a demand $\boldsymbol{d}$, we define $\boldsymbol{d}(p_1,p_2)$ as a demand satisfying, for each user k, $d_k(p_1,p_2)=p_2(d_{p_1^{-1}(k)})$. We can then apply the above bound to demand $\boldsymbol{d}(p_1,p_2)$:

$$R^*(\boldsymbol{d}(p_1,p_2),\boldsymbol{\mathcal{M}})F \geq \sum_{i=1}^{N_{\text{e}}(\boldsymbol{d})}\sum_{j=1}^{F}\mathbb{1}((p_2(d_{u_i}),j)\notin\bigcup_{\ell=1}^{i}\mathcal{M}_{p_1(u_\ell)}). \tag{31}$$

It is easy to verify that by taking the average of (31) over all pairs of (p_1,p_2), only the rates for demands in type $\mathcal{D}_{\boldsymbol{s}(\boldsymbol{d})}$ are counted, and each of them is counted the same number of times due to symmetry. Consequently, we have the following bound:

$$R^*(\boldsymbol{s}(\boldsymbol{d}),\boldsymbol{\mathcal{M}})F = \frac{1}{K!N!}\sum_{p_1\in\mathcal{P}_K}\sum_{p_2\in\mathcal{P}_N}R^*(\boldsymbol{d}(p_1,p_2),\boldsymbol{\mathcal{M}})F \tag{32}$$

$$\geq \frac{1}{K!N!}\sum_{p_1\in\mathcal{P}_K}\sum_{p_2\in\mathcal{P}_N}\sum_{i=1}^{N_{\text{e}}(\boldsymbol{d})}\sum_{j=1}^{F}\mathbb{1}((p_2(d_{u_i},j)\notin\bigcup_{\ell=1}^{i}\mathcal{M}_{p_1(u_\ell)}) \tag{33}$$

$$= \frac{1}{K!N}\sum_{p_1\in\mathcal{P}_K}\sum_{i=1}^{N_{\text{e}}(\boldsymbol{d})}\sum_{m=1}^{N}\sum_{j=1}^{F}\mathbb{1}((m,j)\notin\bigcup_{\ell=1}^{i}\mathcal{M}_{p_1(u_\ell)}). \tag{34}$$

Let a_n be number of bits in the database that are cached by exactly n users, the following equations hold for any i.

$$\frac{1}{K!}\sum_{p_1\in\mathcal{P}_K}\sum_{m=1}^{N}\sum_{j=1}^{F}\mathbb{1}((m,j)\notin\bigcup_{\ell=1}^{i}\mathcal{M}_{p_1(u_\ell)}) = \sum_{n=0}^{K}a_n\frac{\binom{K-n}{i}}{\binom{K}{i}} \tag{35}$$

$$= \sum_{n=0}^{K}a_n\frac{\binom{K-i}{n}}{\binom{K}{n}}. \tag{36}$$

[6]if $M=0$, $\mathcal{R}=\{\{R_K\}_{K\in\mathbb{N}} \mid R_K\geq\min\{N,K\}\}$.

From (34) and (36),

$$R^*(\boldsymbol{s}(\boldsymbol{d}),\boldsymbol{\mathcal{M}})F \geq \frac{1}{N}\sum_{i=1}^{N_{\mathrm{e}}(\boldsymbol{d})}\sum_{n=0}^{K} a_n \frac{\binom{K-i}{n}}{\binom{K}{n}} \tag{37}$$

$$= \frac{1}{N}\sum_{n=0}^{K} a_n \frac{\binom{K}{n+1}-\binom{K-N_{\mathrm{e}}(\boldsymbol{d})}{n+1}}{\binom{K}{n}}. \tag{38}$$

Hence for any $\boldsymbol{s}\in\mathcal{S}$, by arbitrarily picking a demand $\boldsymbol{d}\in\mathcal{D}_{\boldsymbol{s}}$ and applying the above inequality, the following bound holds for any prefetching $\boldsymbol{\mathcal{M}}$:

$$R^*(\boldsymbol{s},\boldsymbol{\mathcal{M}}) \geq \frac{1}{NF}\sum_{n=0}^{K} a_n \frac{\binom{K}{n+1}-\binom{K-N_{\mathrm{e}}(\boldsymbol{s})}{n+1}}{\binom{K}{n}}. \tag{39}$$

Let c_n denote the following sequence

$$c_n = \frac{\binom{K}{n+1}-\binom{K-N_{\mathrm{e}}(\boldsymbol{s})}{n+1}}{\binom{K}{n}}, \tag{40}$$

and we have

$$R^*(\boldsymbol{s},\boldsymbol{\mathcal{M}}) \geq \frac{1}{NF}\sum_{n=0}^{K} a_n c_n. \tag{41}$$

Note that the following equations hold for every prefetching:

$$\sum_{n=0}^{K} a_n = NF, \tag{42}$$

$$\sum_{n=0}^{K} n a_n = NFr, \tag{43}$$

thus from the convexity of lower convex envelopes, we have the following inequality:

$$R^*(\boldsymbol{s},\boldsymbol{\mathcal{M}}) \geq \mathrm{Conv}(c_r). \tag{44}$$

Consequently,

$$\min_{\boldsymbol{\mathcal{M}}} R^*(\boldsymbol{s},\boldsymbol{\mathcal{M}}) \geq \min_{\boldsymbol{\mathcal{M}}} \mathrm{Conv}(c_r) \tag{45}$$

$$= \mathrm{Conv}(c_r) \tag{46}$$

$$= \mathrm{Conv}\left(\frac{\binom{K}{r+1}-\binom{K-N_{\mathrm{e}}(\boldsymbol{s})}{r+1}}{\binom{K}{r}}\right). \tag{47}$$

□

APPENDIX B
MINIMUM PEAK RATE FOR CENTRALIZED CACHING

Consider a caching problem with K users, a database of N files, and a local cache size of M files for each user. We define the rate-memory tradeoff for the peak rate as follows. Similar to the average rate case, for each prefetching $\boldsymbol{\mathcal{M}}$, let $R^*_{\text{worst}}(\boldsymbol{\mathcal{M}})$ denote the peak rate, defined as

$$R^*_{\text{worst}}(\boldsymbol{\mathcal{M}}) = \max_{\boldsymbol{d}} R^*(\boldsymbol{d},\boldsymbol{\mathcal{M}}). \tag{48}$$

We aim to find the minimum peak rate R^*_{worst}, where

$$R^*_{\text{worst}} = \min_{\boldsymbol{\mathcal{M}}} R^*_{\text{worst}}(\boldsymbol{\mathcal{M}}), \tag{49}$$

which is a function of N, K, M, for sufficiently large F.

Now we prove Corollary 1, which completely characterizes the value of R^*_{worst}.

Proof. It is easy to show that the rate stated in Corollary 1 can be exactly achieved using the caching scheme introduced in Section IV. Hence, we focus on proving the optimality of the proposed coding scheme.

Recall the definitions of statistics and types (see section V). Given a prefetching $\boldsymbol{\mathcal{M}}$ and statistics $\boldsymbol{s}$, we define the peak rate within type $\mathcal{D}_{\boldsymbol{s}}$, denoted by $R^*_{\text{worst}}(\boldsymbol{s},\boldsymbol{\mathcal{M}})$, as

$$R^*_{\text{worst}}(\boldsymbol{s},\boldsymbol{\mathcal{M}}) = \max_{\boldsymbol{d}\in\mathcal{D}_{\boldsymbol{s}}} R^*(\boldsymbol{d},\boldsymbol{\mathcal{M}}). \tag{50}$$

Note that

$$R^*_{\text{worst}} = \min_{\boldsymbol{\mathcal{M}}}\max_{\boldsymbol{s}} R^*_{\text{worst}}(\boldsymbol{s},\boldsymbol{\mathcal{M}}) \tag{51}$$

$$\geq \max_{\boldsymbol{s}}\min_{\boldsymbol{\mathcal{M}}} R^*_{\text{worst}}(\boldsymbol{s},\boldsymbol{\mathcal{M}}). \tag{52}$$

Hence, in order to lower bound R^*, it is sufficient to bound the minimum value of $R^*_{\text{worst}}(\boldsymbol{s}, \boldsymbol{\mathcal{M}})$ for each type $\mathcal{D}_{\boldsymbol{s}}$ individually. Using Lemma 2, the following bound holds for each $s \in \mathcal{S}$:

$$\min_{\boldsymbol{\mathcal{M}}} R^*_{\text{worst}}(\boldsymbol{s}, \boldsymbol{\mathcal{M}}) \geq \min_{\boldsymbol{\mathcal{M}}} R^*(\boldsymbol{s}, \boldsymbol{\mathcal{M}}) \tag{53}$$

$$\geq \text{Conv}\left(\frac{\binom{K}{r+1} - \binom{K-N_{\text{e}}(\boldsymbol{s})}{r+1}}{\binom{K}{r}} \right). \tag{54}$$

Consequently,

$$R^*_{\text{worst}} \geq \max_{\boldsymbol{s}} \text{Conv}\left(\frac{\binom{K}{r+1} - \binom{K-N_{\text{e}}(\boldsymbol{s})}{r+1}}{\binom{K}{r}} \right) \tag{55}$$

$$= \text{Conv}\left(\frac{\binom{K}{r+1} - \binom{K-\min\{N,K\}}{r+1}}{\binom{K}{r}} \right). \tag{56}$$

□

Remark 14 (Universal Optimality of Symmetric Batch Prefetching - Peak Rate). Inequality (54) characterizes the minimum peak rate given a type $\mathcal{D}_{\boldsymbol{s}}$, if the prefetching $\boldsymbol{\mathcal{M}}$ can be designed based on $\boldsymbol{s}$. However, for (52) to be tight, the peak rate for each different type has to be minimized on the same prefetching. Surprisingly, such an optimal prefetching exists, an example being the symmetric batch prefetching, according to Section IV. This indicates that the symmetric batch prefetching is also universally optimal for all types in terms of peak rates.

APPENDIX C
PROOF OF THEOREM 2

To completely characterize $\mathcal{R}$, we propose decentralized caching schemes to achieve all points in $\mathcal{R}$. We also prove a matching information-theoretic outer bound of the achievable regions, which implies that none of the points outside $\mathcal{R}$ are achievable.

A. The Optimal Decentralized Caching Scheme

To prove the achievability of all points in $\mathcal{R}$, we consider the following optimal prefetching scheme: all users cache $\frac{MF}{N}$ bits in each file uniformly and independently. This prefetching scheme was originally proposed in [16]. For convenience, we refer to this prefetching scheme as *uniformly random prefetching scheme*.

Given this prefetching scheme, each bit in the database is randomly cached by random subsets of the K users. During the delivery phase, we first greedily categorize the bits based on the number of users that cache the bit, and then for each category, we deliver the corresponding messages in an opportunistic way using the delivery scheme described in Section IV for centralized caching.

Consider any demand $\boldsymbol{d}$, where K users are making requests. For each realization of the prefetching on these K users, we break the bits in the database into $K+1$ sets: For each $j \in \{0, 1, ..., K\}$, let $\mathcal{B}_j$ denote the bits that are cached by exactly j users. For each $\mathcal{B}_j$, with high probability for large F, each file i contains $\binom{K}{j}(\frac{M}{N})^j(1-\frac{M}{N})^{K-j}F + o(F)$ bits, and for any subset $\mathcal{K} \in \{1, ..., K\}$ of size j, a total of $(\frac{M}{N})^j(1-\frac{M}{N})^{K-j}F + o(F)$ bits in file i are exclusively cached by users in $\mathcal{K}$. Note that effectively all bits in $\mathcal{B}_j$ are cached using symmetric batch prefetching, given the fact that the number of bits from any file i that are exclusively cached by users in any set $\mathcal{K}$ are equal. We can apply the delivery scheme proposed in Section IV on $\mathcal{B}_j$, and all requested bits can be delivered with a communication rate of $(\frac{M}{N})^j(1-\frac{M}{N})^{K-j}\left(\binom{K}{j+1} - \binom{K-N_{\text{e}}(\boldsymbol{d})}{j+1}\right)$.

Consequently, by applying the delivery scheme for all $j \in \{0, 1, ..., K\}$, we achieve the communication rate of

$$R_K = \sum_{j=0}^{K} \left(\frac{M}{N}\right)^j \left(1-\frac{M}{N}\right)^{K-j} \left(\binom{K}{j+1} - \binom{K-N_{\text{e}}(\boldsymbol{d})}{j+1}\right) \tag{57}$$

$$= \frac{N-M}{M}\left(1 - \left(1-\frac{M}{N}\right)^{N_{\text{e}}(\boldsymbol{d})}\right) \tag{58}$$

for any demand $\boldsymbol{d}$. Hence, the achieved expected rate under uniform demands given K equals $\mathbb{E}[\frac{N-M}{M}(1-\left(1-\frac{M}{N}\right)^{N_{\text{e}}(\boldsymbol{d})})]$, which dominates all points in $\mathcal{R}$.

B. Converse

To prove a matching outer bound of $\mathcal{R}$, we divide the set of all possible demands into types for each $K \in \mathbb{N}$, and we derive the minimum average rate within each type separately. For any statistics $\boldsymbol{s}$, note that let $R_K^*(\boldsymbol{s}, P_{\boldsymbol{\mathcal{M}}})$ denote the average rate within type $\mathcal{D}_{\boldsymbol{s}}$. Rigorously,

$$R_K^*(\boldsymbol{s}, P_{\boldsymbol{\mathcal{M}}}) = \frac{1}{|\mathcal{D}_{\boldsymbol{s}}|} \sum_{\boldsymbol{d} \in \mathcal{D}_{\boldsymbol{s}}} R_K^*(\boldsymbol{d}, P_{\boldsymbol{\mathcal{M}}}). \tag{59}$$

The following lemma tightly lower bounds the minimum value of $R_K^*(\boldsymbol{s}, P_{\boldsymbol{\mathcal{M}}})$:

Lemma 3. *Consider a decentralized caching problem with N files and a local cache size of M files for each user. For any type $\mathcal{D}_{\boldsymbol{s}}$, where K users are making requests, the minimum value of $R_K^*(\boldsymbol{s}, P_{\boldsymbol{\mathcal{M}}})$ is lower bounded by*

$$\min_{P_{\boldsymbol{\mathcal{M}}}} R_K^*(\boldsymbol{s}, P_{\boldsymbol{\mathcal{M}}}) \geq \frac{M-N}{M} \left(1 - \left(1 - \frac{M}{N}\right)^{N_{\mathrm{e}}(s)}\right). \tag{60}$$

Remark 15. As proved in Appendix C-A, the rate $R_K^*(\boldsymbol{s}, P_{\boldsymbol{\mathcal{M}}})$ for any statistics $\boldsymbol{s}$ and any K can be simultaneously minimized using the uniformly random prefetching scheme. This demonstrates that the uniformly random prefetching scheme is universally optimal for the decentralized caching problem in terms of average rates.

Proof. To prove Lemma 3, we first consider a class of *generalized demands*, where not all users in the caching systems are required to request a file. We define generalized demand $\boldsymbol{d} = (d_1, ..., d_K) \in \{0, 1, ..., N\}^K$, where a nonzero d_k denotes the index of the file requested by k, while $d_k = 0$ indicates that user k is not making a request. We define statistics and their corresponding types in the same way, and we denote the centralized average rate on a generalized type $\mathcal{D}_{\boldsymbol{s}}$ given prefetching $\boldsymbol{\mathcal{M}}$ as $R_K^*(\boldsymbol{s}, \boldsymbol{\mathcal{M}})$.

For a centralized caching problem, we can easily generalize Lemma 2 to the following lemma for the generalized demands:

Lemma 4. *Consider a caching problem with N files, K users, and a local cache size of M files for each user. For any generalized type $\mathcal{D}_{\boldsymbol{s}}$, the minimum value of $R_K^*(\boldsymbol{s}, \boldsymbol{\mathcal{M}})$ is lower bounded by*

$$\min_{\boldsymbol{\mathcal{M}}} R_K^*(\boldsymbol{s}, \boldsymbol{\mathcal{M}}) \geq \mathrm{Conv}\left(\frac{\binom{K}{r+1} - \binom{K-N_{\mathrm{e}}(\boldsymbol{s})}{r+1}}{\binom{K}{r}}\right), \tag{61}$$

where $\mathrm{Conv}(f(r))$ *denotes the lower convex envelope of the following points:* $\{(r, f(r)) \mid r \in \{0, 1, ..., K\}\}$.

The above lemma can be proved exactly the same way as we proved Lemma 2, and the universal optimality of symmetric batch prefetching still holds for the generalized demands.

For a decentralized caching problem, we can also generalize the definition of $R_K^*(\boldsymbol{s}, P_{\boldsymbol{\mathcal{M}}})$ correspondingly. We can easily prove that, when a decentralized caching scheme is used, the expected value of $R_K^*(\boldsymbol{s}, \boldsymbol{\mathcal{M}})$ is no greater than $R_K^*(\boldsymbol{s}, P_{\boldsymbol{\mathcal{M}}})$. Consequently,

$$R_K^*(\boldsymbol{s}, P_{\boldsymbol{\mathcal{M}}}) \geq \mathbb{E}_{\boldsymbol{\mathcal{M}}}[R_K^*(\boldsymbol{s}, \boldsymbol{\mathcal{M}})] \tag{62}$$

$$\geq \mathrm{Conv}\left(\frac{\binom{K}{r+1} - \binom{K-N_{\mathrm{e}}(\boldsymbol{s})}{r+1}}{\binom{K}{r}}\right), \tag{63}$$

for any generalized type $\mathcal{D}_{\boldsymbol{s}}$ and for any $P_{\boldsymbol{\mathcal{M}}}$.

Now we prove that value $R_K^*(\boldsymbol{s}, P_{\boldsymbol{\mathcal{M}}})$ is independent of parameter K given $\boldsymbol{s}$ and $P_{\boldsymbol{\mathcal{M}}}$: Consider a generalized statistic $\boldsymbol{s}$. Let $K_{\boldsymbol{s}} = \sum_i s_i$, which equals the number of active users for demands in $\mathcal{D}_{\boldsymbol{s}}$. For any caching system with $K > K_{\boldsymbol{s}}$ users, and for any subset $\mathcal{K}$ of $K_{\boldsymbol{s}}$ users, let $\mathcal{D}_{\mathcal{K}}$ denote the set of demands in $\mathcal{D}_{\boldsymbol{s}}$ where only users in $\mathcal{K}$ are making requests. Note that $\mathcal{D}_{\boldsymbol{s}}$ equals the union of disjoint sets $\mathcal{D}_{\mathcal{K}}$ for all subsets $\mathcal{K}$ of size $K_{\boldsymbol{s}}$. Thus we have,

$$R_K^*(\boldsymbol{s}, P_{\boldsymbol{\mathcal{M}}}) = \frac{1}{|\mathcal{D}_{\boldsymbol{s}}|} \sum_{\boldsymbol{d} \in \mathcal{D}_{\boldsymbol{s}}} R_K^*(\boldsymbol{d}, P_{\boldsymbol{\mathcal{M}}}) \tag{64}$$

$$= \frac{1}{|\mathcal{D}_{\boldsymbol{s}}|} \sum_{\mathcal{K}:|\mathcal{K}|=K_{\boldsymbol{s}}} \sum_{\boldsymbol{d} \in \mathcal{D}_{\mathcal{K}}} R_K^*(\boldsymbol{d}, P_{\boldsymbol{\mathcal{M}}}) \tag{65}$$

$$= \frac{1}{|\mathcal{D}_{\boldsymbol{s}}|} \sum_{\mathcal{K}:|\mathcal{K}|=K_{\boldsymbol{s}}} |\mathcal{D}_{\mathcal{K}}| R_{K_{\boldsymbol{s}}}^*(\boldsymbol{s}, P_{\boldsymbol{\mathcal{M}}}) \tag{66}$$

$$= R_{K_{\boldsymbol{s}}}^*(\boldsymbol{s}, P_{\boldsymbol{\mathcal{M}}}). \tag{67}$$

Consequently,

$$R_{K_{\boldsymbol{s}}}^*(\boldsymbol{s}, P_{\boldsymbol{\mathcal{M}}}) = \lim_{K \to +\infty} R_K^*(\boldsymbol{s}, P_{\boldsymbol{\mathcal{M}}}) \tag{68}$$

$$\geq \lim_{K \to +\infty} \mathrm{Conv}\left(\frac{\binom{K}{r+1} - \binom{K-N_{\mathrm{e}}(\boldsymbol{s})}{r+1}}{\binom{K}{r}}\right) \tag{69}$$

$$= \frac{M-N}{M}\left(1-\left(1-\frac{M}{N}\right)^{N_{\mathrm{e}}(s)}\right). \tag{70}$$

□

From Lemma 3, we have

$$R_K^*(P_{\boldsymbol{\mathcal{M}}}) = \mathbb{E}_{\boldsymbol{s}}[R_K^*(\boldsymbol{s}, P_{\boldsymbol{\mathcal{M}}})] \tag{71}$$

$$\geq \mathbb{E}_{\boldsymbol{d}}\left[\frac{M-N}{M}\left(1-\left(1-\frac{M}{N}\right)^{N_{\mathrm{e}}(d)}\right)\right], \tag{72}$$

for any $K \in \mathbb{N}$ and for any prefetching scheme $P_{\boldsymbol{\mathcal{M}}}$. Consequently, any vector $\{R_K\}_{K\in\mathbb{N}}$ in $\mathcal{R}$ satisfies

$$R_K^* \geq \min_{P_{\boldsymbol{\mathcal{M}}}} R_K^*(P_{\boldsymbol{\mathcal{M}}}) \tag{73}$$

$$\geq \mathbb{E}_{\boldsymbol{d}}\left[\frac{M-N}{M}\left(1-\left(1-\frac{M}{N}\right)^{N_{\mathrm{e}}(d)}\right)\right], \tag{74}$$

for any $K \in \mathbb{N}$. Hence,

$$\mathcal{R} \subseteq \{\{R_K\}_{K\in\mathbb{N}} \mid R_K \geq \mathbb{E}_{\boldsymbol{d}}[\frac{N-M}{M}(1-(\frac{N-M}{N})^{N_{\mathrm{e}}(\boldsymbol{d})})]\}. \tag{75}$$

Appendix D
Proof of Corollary 2

Proof. It is easy to show that all points in $\mathcal{R}_{\text{worst}}$ can be achieved using the decentralized caching scheme introduced in Appendix C-A. Hence, we focus on proving the optimality of the proposed decentralized caching scheme.

Recall the definitions of statistics and types (see section V). Given a caching system with N files, K users, a prefetching scheme $P_{\boldsymbol{\mathcal{M}}}$, and a statistic $\boldsymbol{s}$, we define the peak rate within type $\mathcal{D}_{\boldsymbol{s}}$, denoted by $R_{K,\text{worst}}^*(\boldsymbol{s}, P_{\boldsymbol{\mathcal{M}}})$, as

$$R_{K,\text{worst}}^*(\boldsymbol{s}, P_{\boldsymbol{\mathcal{M}}}) = \max_{\boldsymbol{d}\in\mathcal{D}_{\boldsymbol{s}}} R^*(\boldsymbol{d}, P_{\boldsymbol{\mathcal{M}}}). \tag{76}$$

Note that any point $\{R_K\}_{K\in\mathbb{N}}$ in $\mathcal{R}_{\text{worst}}$ satisfies

$$R_K \geq \min_{P_{\boldsymbol{\mathcal{M}}}} \max_{\boldsymbol{s}} R_{K,\text{worst}}^*(\boldsymbol{s}, P_{\boldsymbol{\mathcal{M}}}) \tag{77}$$

$$\geq \max_{\boldsymbol{s}} \min_{P_{\boldsymbol{\mathcal{M}}}} R_{K,\text{worst}}^*(\boldsymbol{s}, P_{\boldsymbol{\mathcal{M}}}), \tag{78}$$

for any $K \in \mathbb{N}$. Hence, in order to outer bound $\mathcal{R}_{\text{worst}}$, it is sufficient to bound the minimum value of $R_{K,\text{worst}}^*(\boldsymbol{s}, P_{\boldsymbol{\mathcal{M}}})$ for each type $\mathcal{D}_{\boldsymbol{s}}$ individually. Using Lemma 3, the following bound holds for each $s \in \mathcal{S}$:

$$\min_{P_{\boldsymbol{\mathcal{M}}}} R_{K,\text{worst}}^*(\boldsymbol{s}, P_{\boldsymbol{\mathcal{M}}}) \geq \min_{P_{\boldsymbol{\mathcal{M}}}} R_K^*(\boldsymbol{s}, P_{\boldsymbol{\mathcal{M}}}) \tag{79}$$

$$\geq \frac{M-N}{M}\left(1-\left(1-\frac{M}{N}\right)^{N_{\mathrm{e}}(s)}\right). \tag{80}$$

Hence,

$$\max_{\boldsymbol{s}} \min_{P_{\boldsymbol{\mathcal{M}}}} R_{K,\text{worst}}^*(\boldsymbol{s}, P_{\boldsymbol{\mathcal{M}}}) \geq \max_{\boldsymbol{s}} \frac{M-N}{M}\left(1-\left(1-\frac{M}{N}\right)^{N_{\mathrm{e}}(s)}\right) \tag{81}$$

$$= \frac{M-N}{M}\left(1-\left(1-\frac{M}{N}\right)^{\min N,K}\right). \tag{82}$$

Consequently,

$$\mathcal{R}_{\text{worst}} \subseteq \{\{R_K\}_{K\in\mathbb{N}} \mid R_K \geq \frac{N-M}{M}(1-(\frac{N-M}{N})^{\min\{N,K\}})\}. \tag{83}$$

□

Remark 16. According to the above discussion, the rate $R_{K,\text{worst}}^*(\boldsymbol{s}, P_{\boldsymbol{\mathcal{M}}})$ for any statistics $\boldsymbol{s}$ and any K can be simultaneously minimized using the uniformly random prefetching scheme. This indicates that the uniformly random prefetching scheme is universally optimal for all types in terms of peak rates.

References

[1] D. D. Sleator and R. E. Tarjan, "Amortized efficiency of list update and paging rules," *Communications of the ACM*, vol. 28, no. 2, pp. 202–208, 1985.

[2] L. W. Dowdy and D. V. Foster, "Comparative models of the file assignment problem," *ACM Comput. Surv.*, vol. 14, pp. 287–313, June 1982.

[3] K. C. Almeroth and M. H. Ammar, "The use of multicast delivery to provide a scalable and interactive video-on-demand service," *IEEE J. Sel. Areas Commun.*, vol. 14, pp. 1110–1122, Aug. 1996.

[4] A. Dan, D. Sitaram, and P. Shahabuddin, "Dynamic batching policies for an on-demand video server," *Multimedia Syst.*, vol. 4, pp. 112–121, June 1996.

[5] M. R. Korupolu, C. G. Plaxton, and R. Rajaraman, "Placement algorithms for hierarchical cooperative caching," in *Proc. ACM-SIAM SODA*, pp. 586–595, Jan. 1999.

[6] A. Meyerson, K. Munagala, and S. Plotkin, "Web caching using access statistics," in *Proc. ACM-SIAM SODA*, pp. 354–363, 2001.

[7] I. Baev, R. Rajaraman, and C. Swamy, "Approximation algorithms for data placement problems," *SIAM J. Comput.*, vol. 38, pp. 1411–1429, July 2008.

[8] S. Borst, V. Gupta, and A. Walid, "Distributed caching algorithms for content distribution networks," in *Proc. IEEE INFOCOM*, pp. 1–9, Mar. 2010.

[9] M. A. Maddah-Ali and U. Niesen, "Fundamental limits of caching," *IEEE Trans. Inf. Theory*, vol. 60, pp. 2856–2867, May 2014.

[10] Z. Chen, "Fundamental limits of caching: Improved bounds for small buffer users," *arXiv preprint arXiv:1407.1935*, 2014.

[11] K. Wan, D. Tuninetti, and P. Piantanida, "On caching with more users than files," *arXiv preprint arXiv:1601.06383*, 2016.

[12] S. Sahraei and M. Gastpar, "K users caching two files: An improved achievable rate," in *Proc. CISS*, pp. 620–624, Mar. 2016.

[13] C. Tian and J. Chen, "Caching and delivery via interference elimination," *arXiv preprint arXiv:1604.08600*, 2016.

[14] M. M. Amiri and D. Gunduz, "Fundamental limits of caching: improved delivery rate-cache capacity trade-off," *arXiv preprint arXiv:1604.03888*, 2016.

[15] M. M. Amiri, Q. Yang, and D. Gunduz, "Coded caching for a large number of users," *arXiv preprint arXiv:1605.01993*, 2016.

[16] M. A. Maddah-Ali and U. Niesen, "Decentralized coded caching attains order-optimal memory-rate tradeoff," *IEEE/ACM Trans. Netw.*, vol. 23, pp. 1029–1040, Aug. 2015.

[17] R. Pedarsani, M. A. Maddah-Ali, and U. Niesen, "Online coded caching," in *Proc. IEEE ICC*, pp. 1878–1883, June 2014.

[18] U. Niesen and M. A. Maddah-Ali, "Coded caching with nonuniform demands," in *Proc. IEEE INFOCOM WKSHPS*, pp. 221–226, Apr. 2014.

[19] J. Zhang, X. Lin, and X. Wang, "Coded caching under arbitrary popularity distributions," in *Proc. ITA*, pp. 98–107, Feb. 2015.

[20] M. Ji, A. M. Tulino, J. Llorca, and G. Caire, "Order-optimal rate of caching and coded multicasting with random demands," *arXiv preprint arXiv:1502.03124*, 2015.

[21] J. Hachem, N. Karamchandani, and S. Diggavi, "Multi-level coded caching," in *Proc. IEEE ISIT*, pp. 56–60, 2014.

[22] N. Karamchandani, U. Niesen, M. A. Maddah-Ali, and S. Diggavi, "Hierarchical coded caching," in *Proc. IEEE ISIT*, pp. 2142–2146, June 2014.

[23] J. Hachem, N. Karamchandani, and S. Diggavi, "Effect of number of users in multi-level coded caching," in *Proc. IEEE ISIT*, June 2015.

[24] M. Ji, G. Caire, and A. F. Molisch, "Fundamental limits of caching in wireless D2D networks," *IEEE Trans. Inf. Theory*, vol. 62, pp. 849–869, Feb. 2016.

[25] M. A. Maddah-Ali and U. Niesen, "Cache-aided interference channels," in *Proc. IEEE ISIT*, pp. 809–813, June 2015.

[26] N. Naderializadeh, M. A. Maddah-Ali, and A. S. Avestimehr, "Fundamental limits of cache-aided interference management," in *Proc. IEEE ISIT*, July 2016.

[27] J. Hachem, U. Niesen, and S. Diggavi, "A layered caching architecture for the interference channel," *arXiv preprint arXiv:1605.01668*, 2016.

[28] J. Hachem, U. Niesen, and S. N. Diggavi, "Degrees of freedom of cache-aided wireless interference networks," *arXiv preprint arXiv:1606.03175*, 2016.

[29] C.-Y. Wang, S. H. Lim, and M. Gastpar, "Information-theoretic caching," in *Proc. IEEE ISIT*, pp. 1776–1780, 2015.

[30] C.-Y. Wang, S. H. Lim, and M. Gastpar, "Information-theoretic caching: Sequential coding for computing," *arXiv preprint arXiv:1504.00553*, 2015.

[31] S. H. Lim, C.-Y. Wang, and M. Gastpar, "Information theoretic caching: The multi-user case," *arXiv preprint arXiv:1604.02333*, 2016.

[32] R. Timo and M. Wigger, "Joint cache-channel coding over erasure broadcast channels," in *Proc. IEEE ISIT*, pp. 201–205, 2015.

[33] S. S. Bidokhti, R. Timo, and M. Wigger, "Erasure broadcast networks with receiver caching," in *Proc. IEEE ISIT*, 2016.

[34] S. S. Bidokhti, M. Wigger, and R. Timo, "Noisy broadcast networks with receiver caching," *arXiv preprint arXiv:1605.02317*, 2016.

[35] S. S. Bidokhti, M. Wigger, and R. Timo, "An upper bound on the capacity-memory tradeoff of degraded broadcast channels," in *International Symposium on Turbo Codes & Iterative Information Processing, Brest, France*, 2016.

[36] J. Zhang and P. Elia, "Fundamental limits of cache-aided wireless bc: Interplay of coded-caching and csit feedback," *arXiv preprint arXiv:1511.03961*, 2015.

[37] J. Zhang, F. Engelmann, and P. Elia, "Coded caching for reducing csit-feedback in wireless communications," in *Proc. Allerton Conf.*, pp. 1099–1105, 2015.

[38] S. Li, M. A. Maddah-Ali, Q. Yu, and A. Salman Avestimehr, "A Fundamental Tradeoff between Computation and Communication in Distributed Computing," *ArXiv e-prints*, Apr. 2016.

[39] S. Li, Q. Yu, M. A. Maddah-Ali, and A. S. Avestimehr, "Edge-facilitated wireless distributed computing," in *Proc. IEEE GLOBECOM*, Dec. 2016.

[40] S. Li, Q. Yu, M. A. Maddah-Ali, and A. S. Avestimehr, "A scalable framework for wireless distributed computing," *arXiv preprint arXiv:1608.05743*, 2016.

[41] H. Ghasemi and A. Ramamoorthy, "Improved lower bounds for coded caching," in *Proc. IEEE ISIT*, June 2015.

[42] A. Sengupta and R. Tandon, "Improved approximation of storage-rate tradeoff for caching via new outer bounds," in *Proc. IEEE ISIT*, June 2015.

[43] C.-Y. Wang, S. H. Lim, and M. Gastpar, "A new converse bound for coded caching," *arXiv preprint arXiv:1601.05690*, 2016.

[44] W. Kai, P. Piantanida, and D. Tuninetti, "On the optimality of uncoded cache placement," in *Proc. IEEE ICC*, 2016.

[45] C. Tian, "Symmetry, demand types and outer bounds in caching system," in *Proc. IEEE ISIT*, 2016.